MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

OPERATIONS ANALYSIS OFFICE

SUPPLY READINESS ASSESSMENT:
SOME NOTES ON REPAIR PIPELINE RESPONSES

MARCH 1984

HQ PACIFIC AIR FORCES

84 06 06 010

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO. AD-A141 836	3. RECIPIENT'S CATALOG NUMBER
4. TITLE *(and Subtitle)* SUPPLY READINESS ASSESSMENT: Some Notes on Repair Pipeline Responses		5. TYPE OF REPORT & PERIOD COVERED Final Report
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR*(s)* Robert E. Hiller, Ph.D.		8. CONTRACT OR GRANT NUMBER*(s)*
9. PERFORMING ORGANIZATION NAME AND ADDRESS Hq Pacific Air Forces (OA) Hickam AFB HI 96853		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Same as above		12. REPORT DATE March 1984
		13. NUMBER OF PAGES 19
14. MONITORING AGENCY NAME & ADDRESS*(if different from Controlling Office)*		15. SECURITY CLASS. *(of this report)* UNCLASSIFIED
		15a. DECLASSIFICATION DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT *(of this Report)*

Unlimited

17. DISTRIBUTION STATEMENT *(of the abstract entered in Block 20, if different from Report)*

Unlimited

18. SUPPLEMENTARY NOTES

19. KEY WORDS *(Continue on reverse side if necessary and identify by block number)*

Readiness	Aircraft maintenance	Dyna-METRIC
Logistics	Aircraft repair	Backorders
Supply	Sortie surge	Pipeline Theory
MICAP	NMCS	
Simulation	Sustainability	

20. ABSTRACT *(Continue on reverse side if necessary and identify by block number)*

In previous papers, a Supply Readiness Assessment (SRA) diagram was proposed as a framework within which peacetime MICAP data could be plotted. The MICAP data within the SRA diagram give a quickly comprehensible picture of those items which determine the peacetime supply readiness of the unit. Because the SRA diagram is rooted in and is consistent with the concepts of dynamic pipeline theory for reparable spare parts, that theory can be used to extrapolate the plotted peacetime data into wartime surge predictions. One of

the benefits of the SRA approach is that very useful information can be obtained without detailed and extensive computer programs. That characteristic is just as desirable when the diagram is used predictively. This note provides numerous illustrative examples of repair pipeline responses to various changes in "demand" behavior and "repair time" distributions that describe individual parts. By studying the illustrations and understanding the functional response to changes, it quickly becomes apparent that linear "approximations" give very good extrapolations. Rules of thumb are discussed.

SUPPLY READINESS ASSESSMENT:

SOME NOTES ON REPAIR PIPELINE RESPONSES

Robert E. Hiller, Ph.D.

OPERATIONS ANALYSIS OFFICE
HEADQUARTERS, PACIFIC AIR FORCES
HICKAM AFB HI 96853

MARCH 1984

OVERVIEW

In previous papers, a Supply Readiness Assessment (SRA) diagram was proposed as a framework within which peacetime MICAP data could be plotted. The MICAP data within the SRA diagram give a quickly comprehensible picture of those items which determine the peacetime supply readiness of the unit. Because the SRA diagram is rooted in and is consistent with the concepts of dynamic pipeline theory for reparable spare parts, that theory can be used to extrapolate the plotted peacetime data into wartime surge predictions. One of the benefits of the SRA approach is that very useful information can be obtained without detailed and extensive computer programs. That characteristic is just as desirable when the diagram is used predictively. This note provides numerous illustrative examples of repair pi line responses to various changes in "demand" behavior and "repair ti" distributions that describe individual parts. By studying the illustrations and understanding the functional response to changes, it quickly becomes apparent that linear "approximations" give very good extrapolations. Rules of thumb are discussed.

SUPPLY READINESS ASSESSMENT: Some Notes on Repair Pipeline Responses

I. INTRODUCTION

In some interpretations of the Supply-Readiness Assessment (SRA) diagram, we need to extend peacetime estimates of expected pipeline quantities so we can understand the implications of variations in the underlying parameters. Although the extension can be done by a computer program which solves the convolution integral that yields the exact value of the Poisson parameter, it can also be very readily approximated by rules of thumb. Simple ways of doing this are needed so that the interpretation of the SR diagram can be made by individual observers without having to rely on a computer. We proposed some rules of thumb and set them forth in previously published discussions of the SRA diagram (references 1 and 2).

We have heard informally that some technicians in the field of pipeline modeling consider our proposed rules of thumb "inadequate" or "insufficient." The adequacy or sufficiency of any approximation is always a judgment call, that we don't contest, but we do assert that the call should be made only after the judge is adequately informed. To some extent, then, this rather tediously drawn out set of notes explains the homework lying behind our proposed rules.

We set out here to show how a deterministic repair line, i.e., one with a constant repair time, responds to a step-function increase in daily demands. Although this is the most simple dynamic situation that can be defined, a full understanding of the response is truly fundamental to comprehending more complex situations. It will be shown that deterministic repair lines are basic building blocks. By forming appropriately chosen linear combinations of deterministic lines, one may mathematically treat a great variety of complex demand and repair functions.

The deterministic repair line has additional intrinsic interest because it gives an upper bound on expected pipeline contents compared to all other functions with the same average repair time when the demand intensity is increasing. When the repair distribution function is not known, as is usually the case, the deterministic line gives a well-behaved limiting case.

Of course, if the time-dependent repair and demand distributions are very complicated -- and if we think we know their shapes and parameters well enough to warrant a precise calculation -- then perhaps it is best to call on computer routines to calculate the points. Usually, however, our information is pretty poor and that brings into question the significance of the computer "precision". Even if we do have the requisite data, it is always helpful to have a clear mental picture of what is going on whether a computer is involved or not. Getting a clear mental picture is what this note is all about.

II. BACKGROUND THEORY

We consider a non-homogeneous Poisson "demand" process that generates broken parts which are fed into a repair process. Hillestad and Carrillo (ref 3) have shown that if the repair process is independent of the demand process, then the number of parts in the pipeline will be Poisson distributed with a mean value given by $\lambda(t)$ where

(1) $$\lambda(t) = \int_0^t m(s) \, \overline{F}(s,t) ds \, .$$

$\overline{F}(s,t)$ is the probability that a part going into the repair process at time s is still in repair at time t, depends only on t-s, and

m(s) is the demand intensity function at time s, discussed more fully later.

We presume for our discussion that the reader is familiar with Hillestad and Carrillo's paper, which gives the theoretical basis for the several computer models that Rand has named "Dyna-METRIC," abbreviated DM herein.

We note, in passing, that the central DM calculation is that which provides $\lambda(t)$ for each NSN. That variable, $\lambda(t)$, contains <u>all</u> of the essential results; all the other calculations are only elaborations and interpretations of what the collection of the $\lambda(t)$'s tell us. Before going on, we should emphasize that $\lambda(t)$ <u>is</u> the mean (or expected) quantity in the pipeline. We have referred to it <u>in</u> describing the SR diagram as the EPQ, i.e., the expected pipeline quantity. In essence, the entire dynamic behavior of the pipeline is contained in the time-dependent behavior of the EPQ. Once that is in hand, the story has been told; all that remains is an elaboration of the theme. If we want to understand time-dependent repair processes, then, we need to understand the implications of equation (1); and one way to develop that understanding is to study in detail first some simple functions and then go on to more elaborate ones. We now start that process.

The integral of eqn (1) can be written as a sum:

(2) $$\lambda(t) = \sum_{s=0}^{s=t} m(s)\overline{F}(s,t)$$

if m(s) and $\overline{F}(s,t)$ can be adequately described by discrete functions.

In most applications of the theory, the discrete form is used because it gives sensibly the same results as continuous functions. This schematization of the real world is typical of model building, where, for instance, we ignore the fact that flying occurs only during part of the 24-hour day. In the case at

hand, the discrete summed form, which looks only at daily intervals, can be entirely adequate, especially in the light of our generally poor knowledge of what the functions $\overline{F}(s,t)$ and $m(s)$ are like in the real world.

It is a simple matter to define $\overline{F}(s,t)$ and $m(s)$ for discrete time intervals of one day and form the vector scalar products of eqn (2) to obtain values of $\lambda(t)$ for each value of t. PACAF's "Vector" model (ref 4) computes the convolutions of functions explicitly defined as vectors, whereas DM may use either closed form or numerical integrations.

Whether we use the integral form (eqn 1) or the discrete model (eqn 2) really doesn't matter, although the discrete one may have some advantages for computer calculations. In the remainder of the paper, we will talk about both.

Basically the process calls for us to write down the demand intensity function, $m(s)$, whose time dependency is that of the sortie rate, the other inputs serving merely as scaling factors. Specifically, one may write $m(s)$ = MD/FH x FH/Sortie x Sorties/Day, where MD/FH is mean demands per flying hour, FH/Sortie is flying hours per sortie, and Sorties/Day is obvious. When a unit starts combat flying, it is the sortie rate that dramatically increases although sortie duration may also change. For this note, the specific values are of less interest than the general behavior.

Here, we focus on a step-function form of the demand intensity. It is a commonly used representation of the transition from peacetime to wartime flying rates.

Figure 1. Step Function Form of Demand Intensity
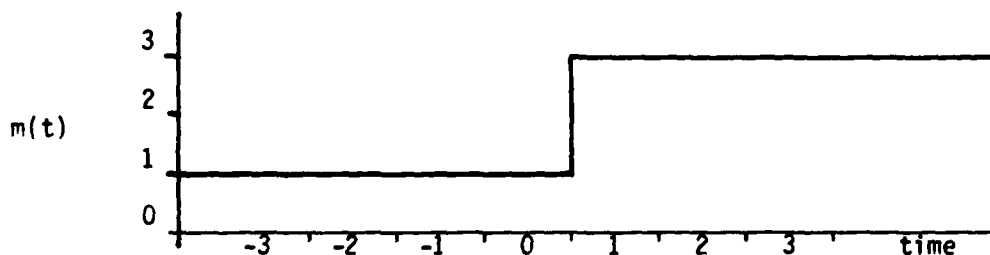
We use the convention that day 0 is the last day of non-surge flying, day 1 is the first day of the surge, and the transition occurs at midnight. We also consider the function to stay constant at 3.0 for future times and to have been 1.0 at all past times. (These, of course, may have any desired absolute values in place of 1.0 and 3.0 which were chosen only for illustration.)

The repair function that represents a "deterministic" or "constant repair time" pipeline is shown in figure 2.

Figure 2. $\overline{F}(t-s)$ for Deterministic Repair Line
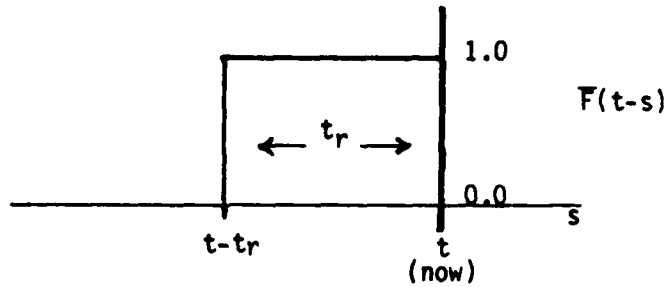
Here, t represents the "present moment," while t_r is the constant repair time. The "probability of a part entering repair" at a time earlier than $t-t_r$ and still being in the line at t is zero, whereas the probability that a part which entered less than t_r ago is sure to be there still, i.e., with probability one.

To evaluate the convolution integral (or the discrete sum), we position the point corresponding to "now" of $\overline{F}(t-s)$ on point t of the intensity function and integrate (or sum) over all non-zero values of the product up to "now," i.e., "t." Then we move $\overline{F}(t-s)$ to a new value of "t" and again integrate (sum), thereby calculating $\lambda(t)$.

In short, the value of the convolution integral is merely the "area" of m(s) which can be seen through (i.e., weighted by) the window $\overline{F}(t-s)$. Since $\overline{F}(t-s)$ has the value of unity everywhere it is non-zero, it acts like an "adder" of the demands for the period t_r when applied to m(s).

Observe that when "t" is in the peacetime region, the adder sees only constant values and $\lambda(t) = t_r m_{ns}$, where m_{ns} is the non-surge value. Similarly, when the surge has been going on longer than t_r, it is clear that $\lambda(t) = t_r m_s$, where m_s is th demand intensity during the surge. During the non-stationary period when $\overline{F}(t-s)$ covers the step portion of m(s), the value of $\lambda(t)$ depends on how much is seen of each portion.

Rather obviously, now, the time dependence of $\lambda(t)$ is that shown below (assume $t_r=5$, $m_{ns}=1.0$, and $m_s=3.0$ for illustration).

Figure 3. Pipeline Contents Response to Step Function Increase in Demands

5
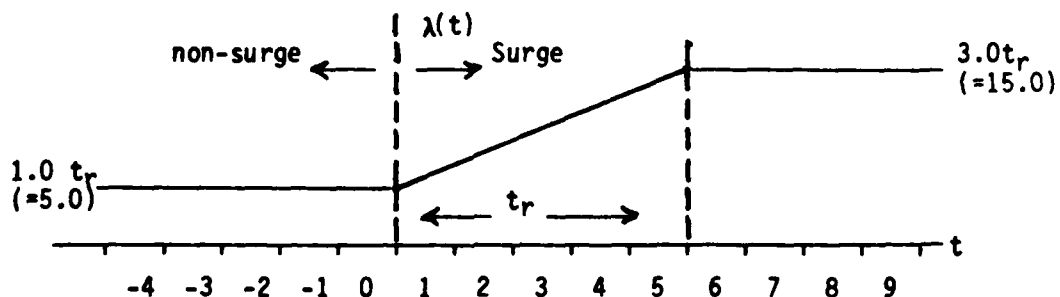
To explore this function, we can apply whatever scaling factors are appropriate, either to m(s) or to the repair time. If the repair time, t_r, is longer, the ramp is longer and vice versa. Outside the ramp, i.e., outside the transient period, the pipeline is "stable" at its stationary-state value.

A very long duration repair line means it will take a long time for the line to stabilize but, conversely, short-time repair lines stabilize very quickly.

We're now in a good position to look at several variations on the repair function, one of which is shown in figure 4. Here the former step function is modified so that a ramp replaces the step. In the discrete form, the functions would look like the "stair step" as shown by the dotted line.

Figure 4
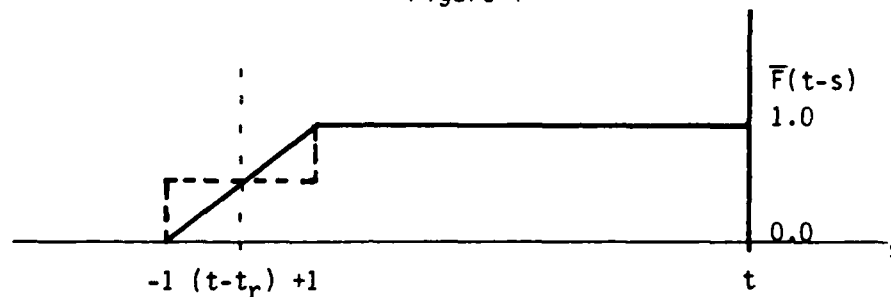
This kind of $\overline{F}(t-s)$ describes a repair process in which the part spends at least a given amount of time in repair and has a uniform probability of getting out of repair during the ramp portion. The average repair time stays the same as before, t_r. Clearly, this "window" gives the same value as the step-function when we're in the regions of stationary demand intensities. (We begin also to see why in the stationary-state Palm's Theorem the pipeline Poisson parameter doesn't depend on the form of the repair distribution, only on average time of repair.) As the index moves across the demand intensity step, the behavior of $\lambda(t)$ is the same as before until $(t-t_r)+1$ arrives at the step, at which point the slope begins decreasing as shown by the curved line for days 5 and 6 of figure 5.

Figure 5

6
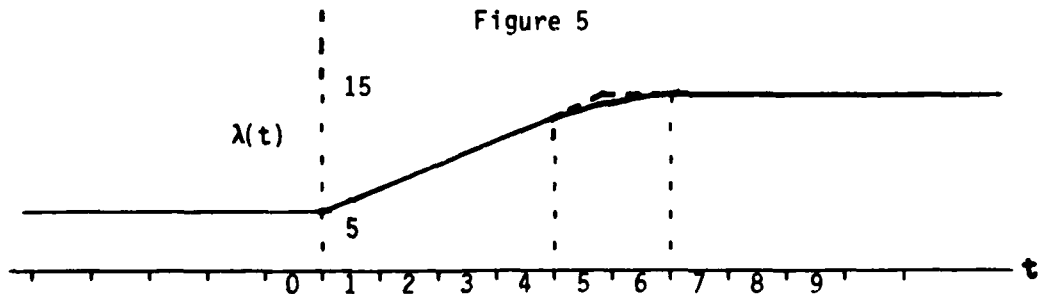
The $\lambda(t)$ transient period between stationary states is now broken into two parts--a ramp and a curve. Obviously, if the period of uniform probability of repair were four days instead of two, the second part would begin after t=3 and last through t=7.

When dealing with discrete functions and a "resolution" unit of one day, the $\overline{F}(t-s)$ of figure 4 can be written as the weighted sum of two step-repair functions, one lasting four days, the other lasting six, with the weights for each being one-half. The convolution integral splits into two equally weighted ones (with appropriate average repair times), and the $\lambda(t)$ response is the weighted average of the 4-day and 6-day step-function responses. Clearly, we can extend the idea and build any kind of "reasonable" $\overline{F}(t-s)$ we wish by creating linear combinations of step-function elements. Thus, the deterministic repair line emerges as a truly fundamental building block, for all complex repair time distributions can be expressed as combinations of them.

In PACAF, a reparable part may go into one of several repair lines, $L(i)$, with probability p_i. We can treat these "parallel" repair lines as a single complex line, or keep them separate if we wish. The theoretical justification for expressing such complex repair distributions by sums of appropriately weighted deterministic repair lines is provided by the theorem that Poisson arrivals with mean m, independently registered into a component pipeline with probability f_i, produce Poisson arrivals at the component pipeline with mean $f_i m$.

Returning to a discussion of $\overline{F}(t-s)$ with shapes like that of figure 4, we can see from the way the response is determined that as long as the repair function ramps are symmetrical about $(t-t_r)$, i.e., the expected repair time doesn't change, the response of $\lambda(t)$ is always equal to or less than that of the $(t-t_r)$ deterministic line. Consequently, the $(t-t_r)$ deterministic line provides an upper bound on expected pipeline quantities for this group of repair functions. This observation, in fact, generalizes to all $\overline{F}(t-s)$ that satisfy the requirements which must be met by the service time distribution: It must start at zero somewhere on the left and rise monotonically to 1.0 at or before the index time "t." The right-most value of 1.0 for $\overline{F}(t-s)$ guarantees that the slope of $\lambda(t)$ starts off at the beginning of the surge with the same value as it does for the $(t-t_r)$ step function, the eventual stationary state value is given by mt_r (Palm's Theorem stationary-state value), and monotonicity takes care of the rest, assuring that the pipeline quantity is never greater than that due to the $(t-t_r)$ step function when t_r is equal to the average repair time. Thus, the $(t-t_r)$ step function provides an upper bound to the transient behavior of all other repair functions having the same average repair time when the average pipeline contents are increasing. Similarly, it provides a lower bound when the average contents are decreasing. For this reason especially, as well as being a basic function for linear expansions, the constant-repair time distribution takes on a "fundamental" character.

Gordon Crawford proved the "bounding-value" property in section IV of his 1981 Rand publication, "Palm's Theorem for Non-Stationary Processes," R-2750-RC, using much more elegant mathematics and notations. Although

PACAF/OA had long before built the "Vector" model (ref 4) on the computational procedures discussed herein, the upper-bound characteristic of the constant repair time distribution had not been fully appreciated until Crawford pointed it out. His paper is, in its entirety, a superb piece of work packed with information and insights. It's worth several careful readings.

A Look at More Complicated Demand Intensities

When a flying unit surges, it does so for only a relatively short time, perhaps up to 8 or 10 days, before it drops back to a lower level of flying activity. The demand intensity function correspondingly drops. Thus, instead of looking just at the step up, we also need to look at the follow-on step down.

Figure 6. Demand Function with an Up-Step and a Down-Step
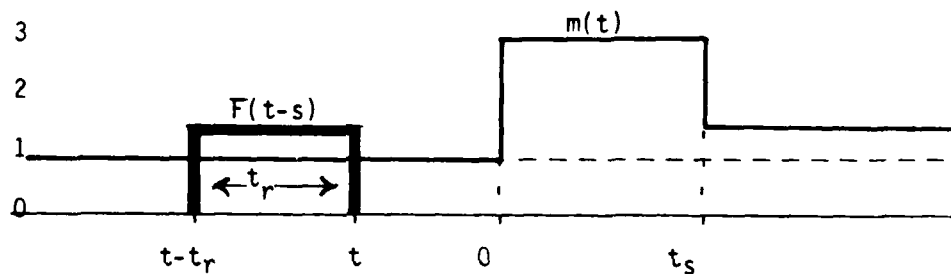
The resulting program could be that shown in figure 6 where t_s is the duration of the surge. We have also marked $\overline{F}(t-s)$ in heavy lines (without worrying about normalization or scaling). With t_s greater than t_r, $\lambda(t)$ increases just as before, stabilizes when $\overline{F}$'s "trailing edge" passes the up-going step, starts down again in a second transient phase when $\overline{F}$'s leading edge passes the down-step, and stabilizes again when $\overline{F}$ is entirely in the right-most stationary-state region. Obviously, the "plateau" of $\lambda(t)$ gets narrower as t_r approaches t_s and is never reached when t_r is greater than t_s.

The behavior is shown in figure 7 where t_s is made successively smaller with respect to t_r. If $\overline{F}$ has a trailing ramp instead of a step, those transition points on $\lambda(t)$ which are affected by the trailing edge would be filleted to a degree determined by the length of the $\overline{F}(t-s)$ ramp, and these are suggested in the figure by the dashed fillets.

The overall behavior shown in figure 7 is exactly what one would intuitively expect when the surge gets short relative to the repair times of the pipeline.

Figure 7. Pipeline Response to 2-Step Demand Intensity Function

Case I: $(t_r < t_s)$ Case II: $(t_r = t_s)$ Case III: $(t_r > t_s)$
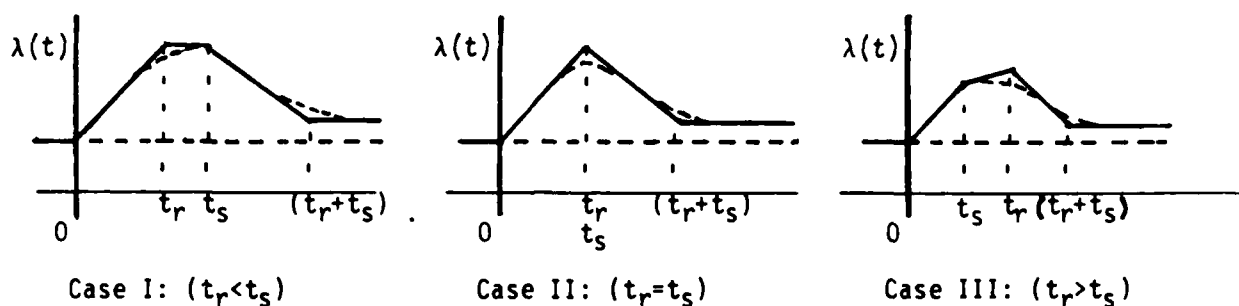
Our exposition so far has treated time-dependent demand functions, but not time-dependent repair functions. We have looked only at those repair functions which depend on time differences, i.e.,

$$F(t,s) = F(t-s).$$

It is quite likely that the repair function will undergo changes at the same time the demand-intensity changes: When the unit begins a flying surge, the maintenance folks will begin to work longer hours, perhaps even be augmented with additional men, so that the average elapsed time to repair a given part could very well decrease at the start of the surge. To round out our story, then, we need to explore the effects of time-dependent changes in the repair function. (Note that the changes must be preordained to occur at points in time, not to occur as a consequence of increasing pipeline contents per se. The repair time distribution must remain statistically independent of the repair line contents.)

The "time-dependent Palm's Theorem" proved in Hillestad and Carrillo's paper covers the general case, but it is easier to make a few points we have in mind by talking in terms of discrete epochs of time. When dealing with non-stationary systems, the theorem can be interpreted as a statement about an ensemble of identical systems, each operating independently but all operating according to the stochastic process described by $m(s)$ and $F(t,s)$. The theorem asserts that the contents of each pipeline in the collection that make up the ensemble will be Poisson distributed with an expectation given by equation (1). To calculate that expectation, one needs know only the "history" of the process, not its future. Thus, we can specify for t_1 what $m_1(s)$ and $F(t_1,s)$ has looked like up 'til t_1 and compute away, labelling the result $\lambda(t_1)$. At another time, t_2, $t_2 > t_1$, we can specify $m_2(s)$ and $F(t_2,s)$ and label the computational result $\lambda(t_2)$. The two calculations are completely unassociated; indeed, they are individually unique. Now, however, if we make $m_1(s)$ and $m_2(s)$ behave in such a way that they are historically consistent with a real-world process, and similarly for $F(t_1,s)$ and $F(t_2,s)$, then we can reasonably say that the sequence $[\lambda(t_1), \lambda(t_2), \ldots, \lambda(t_k)]$ describes the time dependence of the statistics of the ensemble.

9

What does "historical consistency" mean? Wel., with regard to m(s),
there is almost no constraint between "today's" value and "yesterday's" value.
As long as today's history includes verbatim "yesterday's" history, the func-
tion is acceptable. It is essentially a point function over time to which new
points can be freely added. We face a more difficult constraint on $\bar{F}(t,s)$
for, as we have already discussed, the repair line acts as an "adder" of
previous demands. We need to be very careful, then, in deciding whether
$\bar{F}(t_2,s)$ is historically consistent with $\bar{F}(t_1,s)$. We note, again, that
this is not a requirement of the mathematics: We must meet it in order to
claim that our model is in consonance with the world we are modeling. Radical
forms for $\bar{F}(t,s)$ can, for instance, "create" or "destroy," "suck up" or
"return" parts in weird ways.

When the time-dependent Palm's theorem was implemented by RAND in the
string of evolving computer programs they call Dyna-METRIC, the following
treatment was implemented:

The demands input to the peacetime line is cut off at the beginning
of the surge and that line returns repaired parts to the user according to
the peacetime repair distribution. A "wartime" repair line begins receiving
the broken parts due to the surge and they are returned according to the
wartime repair distribution. The total parts in the computed pipeline is the
sum of the parts in the component pipelines. This implementation certainly
guarantees the historical consistency of $\bar{F}(t,s)$ before and after the start
of the surge. Although Dyna-METRIC implements only the one break point (we
believe) the procedure could clearly be repeated again and again to generate
functions more complicated than a pre-surge/post-surge repair distribution.

Let us now look at the response of the pipeline contents to a process
like that above. Figure 8 shows the component pipeline contents.

Figure 8. Component Pipeline Contents for
Pre- and Post-Surge Repair Distributions
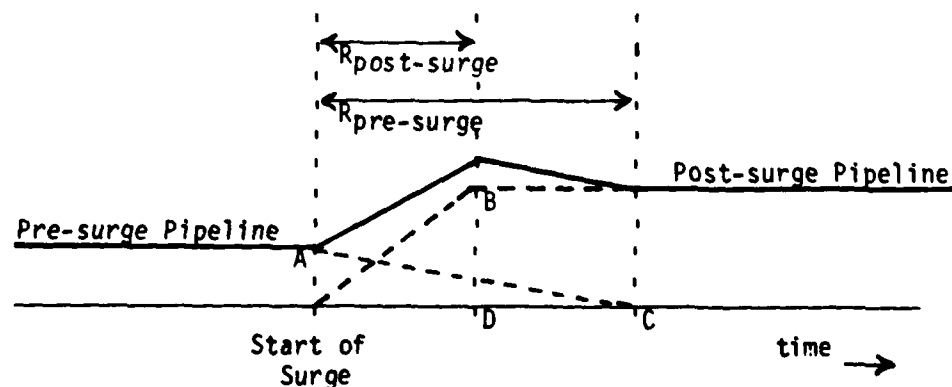
Each component pipeline responds to its specific values of $m(s)$ and $\bar{F}(t,s)$ as shown. Because the "new" post-surge repair time is shorter than the pre-surge repair time, there is a peak in repair line contents. For the values shown ($R_{post-surge} = 0.5\ R_{pre-surge}$; post-surge flying rate = 2 x pre-surge rate), the peak is 25% higher than the long-term steady-state value. If the post-surge repair time were shorter still, the peak would be even higher, and conversely.

Now in point of fact we doubt that the real world will behave the way DM says it will. The stepped-up maintenance capability, whether due to longer hours or extra manpower, could very well work to reduce the peacetime contents very quickly, causing the long ramp, A to C, to become a short one. If it emptied the peacetime contents so that the ramp dropped to zero at D (i.e., parts remaining in the peacetime pipeline are repaired as fast as parts entering during surge), then the solid black line representing total pipeline contents would go directly from A to B and level off without experiencing a peak. If the peacetime contents were repaired even faster than the surge parts (it is possible to rationalize such a case), then the AB ramp would be on the high side, but DM would be even higher.

We hope, at this point, that the reader is receptive to distinctions between the mathematics implicit in the model, the model's congruence to the world it is supposed to represent, and approximations to the mathematics. Of all the many ways to get wrong answers when studying these kinds of problems, the most probable are incorrect data and lack of model congruence to the real world. We will have still more to say about model congruence in the next section but for now will rest on the above example which shows that the way the analyst chooses to represent wartime changes in repair time distributions does have an effect on the calculated peak pipeline quantities. We find it just as reasonable to postulate a model that produces a straight ramp as one which goes above or below it. Similarly, when we do not know either the repair distribution form in peacetime or how it transitions to wartime, it seems silly to worry about the "rounding of the ramps" shown in the illustrative figures.

Our focus in this section has been on showing how a pipeline responds to various simple changes in the functions that drive it. We have used the simplest forms -- steps and ramps -- because (1) they are simple, (2) they are the ones usually implemented when the model is applied, (3) more complex functions can be made up of linear combinations of them, and (4) the determistic repair function is a bounding case. When a user of any of the time-dependent pipeline models (Vector or Dyna-METRIC) has truly grasped the pipeline's behavior we have discussed, he will have gone a long way toward freeing himself from reliance on computer programs. He can use his intellect to explore model reactions and sensitivities with much greater facility than when he's tied cripplingly to sifting computer output data. It was in the hope that more practitioners would start using their own "mental computer" to explore logistics problems that led us to set up the guiding SRA space within which the dynamics are easy to implement.

III. APPLICATION TO SURGE EXTRAPOLATIONS IN THE SR SPACE

The pipeline response $\lambda(t)$, which we have discussed above, is the parameter of the Poisson distribution that describes the pipeline quantity. It is, as previously noted, the EPQ of the SR diagram since the Poisson parameter is the distribution expectation. Based on our understanding of the way the pipeline responds, as developed above, we can now make quite credible extrapolations of peacetime "stationary-state" values. We have elsewhere described the simple rules-of-thumb that can be used. (Section V of PACAF/OA paper, "Readiness Assessment and Cannibalization," (ref 1); and section IV of PACAF/OA paper, "Reparable Item Supply Readiness Assessment Using MICAP Data," Second Draft, ref 2)). The first reference is quoted in the following text, and we have just covered the justification for the rules cited in it.

"V. SR SPACE DYNAMICS

"If delicately shaped calculations are needed and if all the assumed conditions are met, Vector or Dyna-METRIC can tell us how the SR-plot is mapped into another SR-plot at a different time. We have just illustrated it in the previous section. There may well be occasion, however, when such delicacy is not needed and when approximations will suffice. The occasion may arise because of an urgency to get answers, because of economy, because the basic input data are too shaky to warrant precision, and so on. Or it may just be useful to have an idea of how an SR plot is apt to change: such a "feel" is at least part of understanding the SR space. We will later rely on the intrinsic behavior to discuss a few large issues, so it is worthwhile dallying a bit to develop the "feel."

"First, we grasp the idea that the SR space is determined by the repair pipeline distribution. We have, so far, restricted ourselves to a Poisson distribution.

"Second, it eases matters to think at first of stationary conditions. Peacetime training flying is a stationary state.* Even a surge becomes a stationary state if, in our imagination, we let it go on long enough. We can get some very useful benchmarks that way, simply by applying Palm's theorem, the one which preceded the Dynamic Palm's Theorem. In that case, the calculation of the EPQ is very simple: it is just the NSN's daily demand rate multiplied by its mean repair time and, moreover, does not depend on the form of the repair time distribution. When the sortie rate changes and after the pipeline reaches the new steady state, the new EPQ_f will be the product of the new daily demand rate and the new mean repair time. Most simply, if the mean repair time doesn't change, then

---

* Note added: Recent evidence suggests that this widely used assumption may be far from the truth.

13

$$EPQ_f = \frac{(Sortie\ Rate)_f}{(Sortie\ Rate)_i}\ EPQ_i$$

and if it does, the scaling is obvious.

"About now, we should say how long is needed for the new steady state to be reached. For a constant time-of-repair having a value of, say, 10 days, it takes exactly 10 days to reach the new equilibrium. The repair line "memory" is only 10 days, for by then all NSNs from the previous steady state have been repaired and are gone. If the repair time is a constant 30 days, it takes 30 days for the pipeline to forget the past. When the pipeline doesn't remember the past, steady state has been reached. With repair times which are not quite so constant, the mean time of repair is still a generally good approximation to the duration of the transient period."

The text which appears in the second reference says the same thing, although at greater length. Both are precisely correct extrapolation techniques for (1) a step-function surge and (2) a constant service-time distribution, as long as (3) we don't reach the drop-off or step-down point of the surge, i.e., its end. To the extent "approximation" is mentioned, we had in mind repair distributions with ramps, either linear or curvy, on the trailing edge. For those, we have seen that the constant time-of-repair yields an upper bound or "pessimistic" value for EPQ(t) during an up-transient. At the same time, the diagrams show it is not very pessimistic for common values of surge duration and average repair times and that those differences from step-function form don't have much influence. When the repair time distribution changes, then we have the problem of how to guarantee "historical" consistency. DM guarantees it one way, but that way may not accurately portray the likely real-world response. In that event, our approximation may be better than DM. This is a question of how well the model approximates the world, not of how to do the calculation.

Crawford (ref 5) has pointed out -- indeed, we also have been at some pains in this note to make obvious -- that the stationary behavior of the pipeline depends only on the average repair time, i.e., is independent of the shape of the distribution, but that the transient behavior does depend on the shape. Crawford asserts, and we agree, that the second moment of the distribution catches most of the shape-factor effects. Instead of paraphrasing further, we quote from page 27 of Rand's R-2750-RC:

"LACK OF SENSITIVITY TO THE CHOICE OF REPAIR TIME DISTRIBUTION

"In the steady-state calculation less information is needed about repair times. Palm's Theorem says that we do not need to know the distribution of repair times, only the mean repair time.

It was mentioned that the dynamic spares calculation is sensitive to the size of the averaging window--that is, to the mean repair time. The steady-state calculation is no less sensitive to changes in the mean repair time.

"The difference between the steady state and dynamic models is in the sensitivity to the higher moments. The steady-state calculation is insensitive to (in fact independent of) the higher moments. The degree of smoothing in the dynamic calculation is influenced by the size of the tail of the repair time distribution. In Figs. 4, 6, and 7, I have plotted the mean of the number of units in the repair pipeline, by day, using the same flying schedule but three very different repair time distributions having the same mean and the same variance. From the point of view of requirements calculations, the curves are nearly the same. The peaks occur at the same time, and the up-slope parts of the curves that dictate when the extra spares may be needed are almost identical. The large differences in the tails of the distributions cause differences in the curves in the days after the peak, but they would not affect a requirements calculation. For comparison the more peaked $l_0(t)$ corresponding to constant 10 day repair times has been plotted in fig. 8.

"Such examples lead us to believe that if it is possible to estimate the mean and variance of the repair time distribution, it is also possible to use just about any distribution having these moments and the results will be the same."

For our studies of reparable parts within PACAF, we postulate a repair pipeline which is a composite of three independent pipelines: a base repair pipe, a Pacific Logistics Support Center (PLSC) repair pipe, and a depot repair pipe. Each NSN has its own set of probabilities for finding itself in one or the other pipes and its own set of average repair times for each pipe. There was a time when we worried about the shape of the distribution. One of the reasons we built the Vector Model was to explore those effects. We quickly found that nothing very important happened for reasonable variations in service-time distribution shapes. Since the early explorations, we have been content to use the bounding assumption of constant repair times. This composite catches very well the three first moments (i.e., the two RCTs and the depot OS/T) which are dominant both in stationary state and transient periods. In principle, one might want in very precise work to estimate the variance, but we know of no one attempting a large scale application of time-dependent pipeline theory who is trying to do so. Most everyone, it seems, is willing to forego the (loosely speaking) second-order effects of the variance in order to achieve the economies of not having to estimate it.

Thus, in the overall context of present-day modeling of reparable part pipelines, i.e., as they are driven by the peacetime and wartime flying of real world units, a deterministic repair line is a completely acceptable repair time distribution. No analyst working these problems (in large-scale applications, at least) really has collected data on the form of the repair time distribution or estimated the variance for each separate NSN. That is

a minimum precondition to using more complicated distributions than that of the deterministic pipeline. Even more to the point, the repair cycle times themselves are individually not very well known for the whole collection of NSNs that make up the raw material of the model's application. We suspect that few analysts are searching for the second moments simply because they already understand that the more important first moments are only poorly known. Being wrong on the first moment of each NSN's distribution appears to be a far more grievous source of error than the very minor kinds of approximations implicit in the rules of thumb we proposed.

The "linear extrapolation" which seems to have offended some readers now stands exposed as a precise statement of the pipeline response to the most useful, and probably most used, form of the repair time distribution. To the extent it is an approximation to the real world, it is a far less serious and more innocuous approximation than those already imbedded in the "schematic-pragmatic" models of a complicated real world.

We have never asserted or implied a requirement to carry over to war-time surges the peacetime parameters that describe the demand and repair distributions. Nor have we presupposed those values: they are fully at the choice of the user. We believe, in fact, that wartime values will and should be different from peacetime ones. Our proposed rules say nothing that limits the way they change. Whatever their values are, one may still conceive of a long-term steady-state which could be reached and the time needed to reach it. Because EPQ_f is specifically defined as a steady-state value, it may be obtained from proportional relations of the final and initial demand and repair time values. We then use a ramp to show how the daily EPQs progress toward EPQ_f from day to day during the surge. The overall approximation is a good one which should be entirely adequate for the duration of the surge. Indeed it is far better than required for mere discussion of the effects of functional changes in the parameters . . . and that, after all, was our stated purpose.

REFERENCES

1. Readiness Assessment and Cannibalization, HQ PACAF/OA, March 1983

2. Reparable Item Supply-Readiness Assessment Using MICAP Data,
   HQ PACAF/OA, 28 June 1983

3. Models and Techniques for Recoverable Item Stockage When Demand and
   the Repair Process Are Nonstationary--Part I: Performance Measure-
   ment, N-1482-AF, The Rand Corporation, May 1980

4. VECTOR User's Manual, HQ PACAF/OA Technical Note 83-1, 11 March 1983

5. Palm's Theorem for Nonstationary Processes, R-2750-RC, The Rand
   Corporation, October 1981

DISTRIBUTION

HQ USAF Wash DC 20330		LMSC/LD (Lt Col Tripp) Wright-Patterson AFB OH 45433	1
SAG	2		
LE	1	LOC/CC Wright-Patterson AFB OH 45433	2
HQ AFLC Wright-Patterson AFB OH 45433	2	OO-ALC Hill AFB UT 84406	
XR		MM	1
		XR	1
HQ MAC Scott AFB IL 62225			
LG	1	OC-ALC Tinker AFB OK 73145	
XP	1	MM	1
		XR	1
HQ PACAF Hickam AFB HI 96853			
LG	1	SM-ALC McClellan AFB CA 95652	
OA	20	MM	1
		XR	1
HQ SAC/NR Offutt AFB NE 68113	1		
		SA-ALC Kelly AFB TX 78241	
HQ TAC Langley AFB VA 23665		MM	1
LG	2	XR	1
XPS	1		
		WR-ALC Robins AFB GA 31098	
HQ USAFE APO New York 09012		MM	1
DOA	1	XR	1
LG	1		
		5AF/CC APO San Francisco 96328	2
USAFA/DFMS (Maj R. Rue) USAF Academy CO 80840	1		
		13AF/CC APO San Francisco 96274	2
AFDSDC Gunter AFS AL 36114			
CC	1	313AD/CC APO San Francisco 96239	2
LGS	1		
LGX	1	314AD/CC APO San Francisco 96570	2
AFHRL/LG (Mr F. Maher) Wright-Patterson AFB OH 45433	1		
		326AD/CC Wheeler AFB HI 96854	2
AFIT/LSM Wright-Patterson AFB OH 45433	3		
AFLMC/LG Gunter AFS AL 36114	3		

```
3TFW                                    DLSIE/ALMC                      1
APO San Francisco 96274                 Fort Lee VA 23801
CC                              1
DCM                             1
RSM                             1

8TFW
APO San Francisco 96264
CC                              1
DCM                             1
RSM                             1

18TFW
APO San Francisco 96239
CC                              1
DCM                             1
RSM                             1

51TFW
APO San Francisco 96570
CC                              1
DCM                             2
RSM                             2

JUSMAG-K/MKDA                   1
APO San Francisco 96301

PLSC/CC                         1
APO San Francisco 96239

Dynamics Research Corp          1
Attn:  Mr Limpert
60 Concord Street
Wilmington MA 01887

The Rand Corporation            2
Attn:  Dr G. Crawford
1700 Main Street
Santa Monica CA 90406

Synergy, Inc                    1
2337 18th Street, N.W.
Wash DC 20009

Logistics Management Institute  1
4701 Sangamore Road
P. O. Box 9489
Wash DC   20016

DTIC                            12
Cameron Station
Alexandria VA 22314
```